arXiv:1605.02824v1 [cs.DB] 10 May 2016

RORS: Enhanced Rule-based OWL Reasoning on Spark

Zhihui Liu, Zhiyong Feng, Xiaowang Zhang, Xin Wang, and Guozheng Rao

School of Computer Science and Technology, Tianjin University, China
Tianjin Key Laboratory of Cognitive Computing and Application, Tianjin, China
{zhihuiliu,zyfeng,xiaowangzhang,wangx,rgz}@tju.edu.cn

Abstract. The rule-based OWL reasoning is to compute the deductive closure of an ontology by applying RDF/RDFS and OWL entailment rules. The performance of the rule-based OWL reasoning is often sensitive to the rule execution order. In this paper, we present an approach to enhancing the performance of the rule-based OWL reasoning on Spark based on a locally optimal executable strategy. Firstly, we divide all rules (27 in total) into four main classes, namely, *SPO rules* (5 rules), *type rules* (7 rules), *sameAs rules* (7 rules), and *schema rules* (8 rules) since, as we investigated, those triples corresponding to the first three classes of rules are overwhelming (e.g., over 99% in the LUBM dataset) in our practical world. Secondly, based on the interdependence among those entailment rules in each class, we pick out an optimal rule executable order of each class and then combine them into a new rule execution order of all rules. Finally, we implement the new rule execution order on Spark in a prototype called RORS. The experimental results show that the running time of RORS is improved by about 30% as compared to Kim & Park's algorithm (2015) using the LUBM200 (27.6 million triples).

1 Introduction

The Web Ontology Language [7] (OWL) is a semantic web language designed to represent rich and complex knowledge about things, groups of things, and relations between things. There are mainly two modeling paradigms for the semantic web. The first paradigm is based on the notion of the classical logics, such as the description logics algorithms [2] on which the OWL is based. The other paradigm is based on the Datalog paradigm. A subset of the OWL semantics is transformed into rules that are used by a rule engine in order to infer implicit knowledge. This paper focuses on the second paradigm, that is, a rule-based OWL reasoning using OWL-Horst [3] rules.

Owing to the explosion of the semantic data, the number of RDF triples in large public knowledge bases, e.g., DBpedia has increased to billions [4]. Therefore, to improve the performance of OWL reasoning becomes a core problem. The traditional single-node approaches are no longer viable for such large scale data. Some existing ontology reasoning systems are based on MapReduce framework [11,17,18,21]. The OWL reasoning in [21] and [11] perform reasoning over

MapReduce with rule execution mechanism. However, the MapReduce-based approaches are not very efficient due to the data communication between memory and disk. To further improve the performance of reasoning, some researchers have implemented the OWL reasoning on Spark, which is an in-memory and distributed cluster computing framework [22].

Recently, Cichlid[4] has greatly improved the performance of OWL reasoning on Spark as compared to the state-of-the-art distributed reasoning systems, but it only considers parts of OWL rules and does't analyze the interdependence of rules. Reasoning based on OWL-Horst rules can infer many more implicit information. And different rule execution strategy will influence the reasoning performance. For instance, let S be the set of triples of an ontology, where S={$\langle A$ *subClassOf* $B\rangle$, $\langle B$ *subClassOf* $C\rangle$}. R1 and O11 are the two rules of OWL-Horst, where $R1$={$\langle C$ *rdfs:subClassOf* $C1\rangle$, $\langle C1$ *rdfs:subClassOf* $C2\rangle \Rightarrow \langle C$ *rdfs:subClassOf* $C2\rangle$} and $O11$={$\langle V$ *owl:equivalentClass* $W\rangle \Rightarrow \langle V$ *rdfs:subClassOf* $W\rangle$ }. By implementing the R1 entailment rule for subclass closure, we will get that S={$\langle A$ *subClassOf* $B\rangle$, $\langle B$ *subClassOf* $C\rangle$, $\langle A$ *subClassOf* $C\rangle$}. If the O12 entailment rule for equivalent class is executed before the R1, S will contain more new triples and reduce the iterative operation. Therefore, it is desired to optimize reasoning by adjust the rule order. Although Kim & Park [6] has also implemented parallel reasoning algorithms with an executable rule order, but lacked the evidence to prove that the strategy is optimal.

To find the optimal executable strategy, we use the depth-first algorithm to get all possible executable strategies, which are based on the dependency of rules. There are 259367372 possible strategies among the 27 rules in Table 1. Due to the very large number of strategies, it is challenge to find the optimal strategy by test every strategy.

In this paper, we present an approach to enhancing the performance of the rule-based OWL reasoning based on a locally optimal executable strategy and implement the new rule execution strategy on Spark in a prototype called RORS. The major contributions and novelties of our work are summarized as follows:

- We analyze the characteristic of dataset and divide the dataset into three classes: SPO triples, sameAs triples and type triples with analysing the proportion of the three classes respectively in the dataset. According to the data partition, we divide the OWL-Horst rules into four classes.
- We respectively analyze the rule interdependence of each class, and find the optimal executable strategies.
- Based on the locally optimal strategies, we pick out an optimal rule execution order of each class and then combine them into a new rule execution strategy of all rules and implement the new rule execution strategy on Spark.

The rest of this paper is organized as follows. Section 2 gives an brief introduce to preliminary knowledge about OWL and Spark. Section 3 presents our locally optimal strategies. Section 4 implements our proposed strategy on Spark and Section 5 evaluates the experiment on the LUBM dataset. In Section 6, we discuss related works and summarize this paper.

2 Preliminaries

In this section, we briefly recall the ontology language OWL and the framework Spark, largely following the excellent expositions [7,22].

OWL An ontology is a formal naming and definition of the types, properties, and interrelationships of the entities that really or fundamentally exist for a particular domain of discourse. Ontology is part of the W3C standards stacks for the semantic web. The language OWL [7] is a family of knowledge representation languages for authoring ontologies. There are three variants of OWL, with different levels of expressiveness. These are OWL Lite, OWL DL and OWL Full (ordered by increasing expressiveness). Each of these sublanguages is a syntactic extension of its simpler predecessor. OWL DL is designed to preserve some compatibility with RDF Schema (or RDFS). However, OWL Full is undecidable, so no reasoning software is able to perform complete reasoning for it. OWL DL designed to provide the maximum expressiveness possible while retaining computational completeness, decidability, and the availability of practical reasoning algorithms. OWL Lite was originally intended to support those users primarily needing a classification hierarchy and simple constraints. The three languages are one subset of the other.

Spark: Distributed computing framework Spark [22] is an open source cluster computing framework, which is developed at the University of California, Berkeley's AMPLab. One of the main features is the in-memory parallel computing model which all data will be loaded into the memory. Spark provides programmers with an application programming interface centered on a data structure called the resilient distributed dataset [23] (RDD), a read-only multiset of data items distributed over a cluster of machines, that is maintained in a fault-tolerant way. Each RDD will be divided into multiple partitions that exist on different computing nodes. And it provides a variety of operations to transform one RDD into another RDD. There are two kinds of operations. Transformations are lazy operations that define a new RDD (e.g., *map*, *filter*, and *join*), while actions launch a computation to return a value to the program or write data to external storage, such as *collect*, *count*, *saveAsTextFile*, etc. RDD achieves fault tolerance through a notion of lineage based on logging the transformations, if a partition of an RDD is lost, the RDD has enough information about how it was derived from other RDDs to recompute just that partition. More details about Spark please see the official web site `http://spark.apache.org/`.

3 Locally optimal strategy

In this section, we propose a locally optimal strategy based on the dependency among the rules.

Table 1: OWL-Horst rules

Rule ID	Condition	Consequence
R1	c rdfs:subClassOf c1 c1 rdfs:subClassOf c2	c rdfs:subClassOf c2
R2	p rdfs:subPropertyOf p1 p1 rdfs:subPropertyOf p2	p rdfs:subPropertyOf p2
R3	s p o , p rdfs:subPropertyOf p1	s p1 o
R4	s rdfs:domain x , u s y	u rdf:type x
R5	p rdfs:range o , s p v	v rdf:type o
R6	c rdfs:subClassOf c1 , v rdf:type c	v rdf:type c1
O1	p rdf:type owl:FunctionalProperty u p v , u p w	v owl:sameAs w
O2	p rdf:type owl:InverseFunctionalProperty v p u , w p u	v owl:sameAs w
O3	p rdf:type owl:SymmetricProperty v p u	u p v
O4	p rdf:type owl:TransitiveProperty u p w , w p v	u p v
O5	v owl:sameAs w	w owl:sameAs v
O6	v owl:sameAs w w owl:sameAs u	v owl:sameAs u
O7a	p owl:inverseOf q, v p w	w q v
O7b	p owl:inverseOf q, v q w	w p v
O8	v rdf:type owl:Class v owl:sameAs w	v rdfs:subClassOf w
O9	p rdf:type owl:Property p owl:sameAS q	p rdf:subPropertyOf q
O10	u p v, u owl:sameAs x v owl:sameAs y	x p y
O11a	v owl:equivalentClass w	v rdfs:subClassOf w
O11b	v owl:equivalentClass w	w rdfs:subClassOf v
O11c	v owl:subClassOf w w owl:subClassOf v	v rdfs:equivalentClass w
O12a	v owl:equivalentProperty w	v rdfs:subPropertyOf w
O12b	v owl:equivalentProperty w	w rdfs:subPropertyOf v
O12c	v owl:subPropertyOf w w owl:subPropertyOf v	v rdfs:equivalentProperty w
O13	v owl:hasValue w v owl:onProperty p, u p v	u rdf:type v
O14	v owl:hasValue w v owl:onProperty p u rdf:type v	u p v
O15	v owl:someValuesFrom w v owl:onProperty p u p x, x rdf:type w	u rdf:type v
O16	v owl:allValuesFrom w v owl:onProperty p u rdf:type v, u p x	x rdf:type w

3.1 Dependance of rules

OWL-Horst [3] has a powerful expression and reasoning ability. In this paper, the OWL reasoning is based on OWL-Horst rules shown in Table 1. There exists interdependency among the rules. For example, the output of O13 can be the condition of O14, whereas O14 affects O3, O4, O7a, O7b and R3. And more than one rule can be the input of another rule. The output of O3, O4, and R3 affects O13, O15 and O16. If the input to a rule R_i depends on the output of another rule R_j, then the rule R_i is dependent on R_j. The rule R_j should be executed before the rule R_i. A rule dependency graph can be constructed based on the interdependency among the rules. Each vertex represents a rule and an outgoing edge between vertex v_i and v_j represents the dependency of vertex v_i on v_j. The dependency graph of SPO rules is shown in the Figure 1. Based on the dependency graph, there are 259367372 executable strategies among the 27 rules that adjacent rules satisfy dependency, and each rule is executed only once. And the longest strategies contain 22 rules that cannot fully sequentially execute all rules. The number is so huge that it is challenge to test every strategy. It is known that reasoning has close relationship with the characteristic of datasets. Therefore, we prefer to make a analysis of LUBM dataset.

3.2 Locally optimal strategy

LUBM [5] is a widely used standard benchmark for evaluating the performance of ontology reasoning. We can use data generator to generate dataset with different size. Data generator is carried out by the Univ-Bench Artificial data generator (UBA), a tool developed for the benchmark. We firstly divide the dataset into three classes as follows:

- Triples whose predicate is *rdf:type*. We call this class as type.
- Triples whose predicate is *owl:sameAs*. We call this class as *sameAs* .
- The remainder is classified as a class. We call this class as SPO.

The UBA uses specific rules to generate data so that all datasets generated by it have same data characteristic. So we use the dataset LUBM-50 as the sample and analyze the proportion of each class. The result is listed in Table 2.

Table 2: The proportion of each type in LUBM-50

Dataset	type	sameAs	SPO
LUBM-50	20.055%	0	79.945%

From the Table 2, we can see that the SPO triples of LUBM_50 account for absolute proportion, about 80% of the total. The type triples are in the second place, but the number of *sameAs* triples is zero. Therefore, the OWL reasoning should focus on the reasoning of SPO triples and type triples. The *sameAs* triples are simply handled. Based on the statistic, we divide the OWL-Horst rules into four classes as follow:

- Rules whose condition or consequence has triples of type class, including R4, R5, R6, O13, O14, O15, O16. These rules are used to infer implicit type data. The dependency graph of type rules is shown in Figure 2.

– Rules whose condition or consequence has triples of *sameAs* class, including O1, O2, O5, O6, O8, O9, O10. There are certain rules for ontology merging(O8 and O9) [6], so we also exclude these rules from our reasoner. The dependency graph of sameAs rules is shown in Figure 3.
– Rules whose condition or consequence has triples of SPO class, including R3, O3, O4, O7a, O7b. The dependency graph of SPO rules is shown in Figure 1, which the rule O7a and O7b are classified as O7.
– The remainder is classified as a class, including R1, R2, O11a, O11b, O11c, O12a, O12b, O12c. The dependency graph of schema rules is shown in Figure 4.

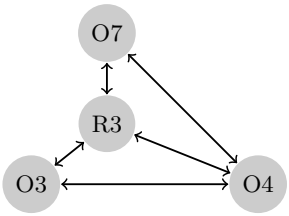

Fig. 1: SPO rules dependency graph

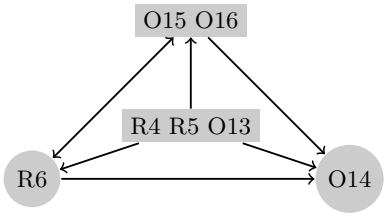

Fig. 2: type rules dependency graph

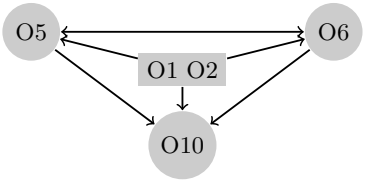

Fig. 3: sameAs rules dependency graph

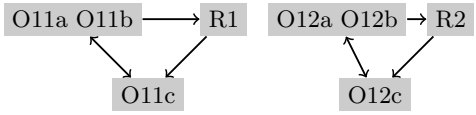

Fig. 4: schema rules dependency graph

For each class, based on the dependency of rules, we use *depth-first-search* (DFS) algorithm to find all possible executable orders among the rules and acquire the optimal executable orders. Through the experiment, there is a large number of orders for each class, but we should chose the longest orders that contain rules as many as possible. The optimal executable orders for each class are listed in Table 3.

Table 3: The optimal strategies of each class

SPO rules	type rules
	R4 → R6 → O14 → O13 → O15 → O16
O3 → R3 → O7 → O4	R4 → R6 → O14 → O13 → O16 → O15
O7 → R3 → O3 → O4	R5 → R6 → O14 → O13 → O15 → O16
	R5 → R6 → O14 → O13 → O16 → O15
schema rules	sameAs rules
O11a, O11b → R1 → O11c	O1 → O10 → O2 → O6 → O5
O12a, O12b → R2 → O12c	O2 → O10 → O1 → O6 → O5

The symmetric rule O3 and the *inverseOf* rule O7 are not affected by each other because a property can not simultaneously involve these two feathers [6]. So O7 and O3 don't have dependence. The transitive rule O4 is used to compute the SPO triples closure and should be applied at last in SPO rules before generating more new triples. The second part is connected by type triples. We find that each path lacks R4 or R5 because there is no type triples in their condition, so they should be classified to same kind of rule. The R4 and R5 can be executed in any order. The third part works with schema triples. Because the O11a and O11b produce similar triples with the only different order of subject and object, so there is no order between O11a and O11b. In the condition of O1 and O2, owing to the different property of the predict, the paths are equivalent when exchanging the order between O1 and O2. The O5 does not produce new triples actually and should be placed at last.

4 Distributed rule-based OWL reasoning on Spark

There are two kinds of triples in OWL reasoning. The schema triples are OWL rule schema which provides basic elements for the description of ontology while the instance triples are actual statement in ontology, including SPO triples, type triples and sameAs triples. After proposing the locally optimal orders for each class, we design the overall strategy of the OWL reasoning shown in Figure 5.

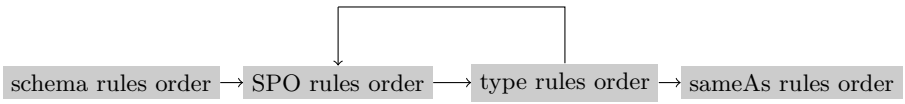

Fig. 5: The reasoning strategy

The reasoning of instance triples will use the output of schema triples, so the schema reasoning should be executed firstly. There are several orders for each class of instance triples. We can chose any one order and then combine an new executable strategy. In this paper, we select first order for each class to perform reasoning. The general workflow of the parallel OWL reasoning is described in Algorithm 1.

Optimize join operation It is inevitable that there exists a lot of join operations on triples, especially the multiple triples in the condition, such as O14. The join operation can greatly influence the performance of reasoning. Generally, the schema triples of an ontology is small even if the ontology is very large [6]. Spark provides the broadcast variables which can transform local data to all available computing nodes and keep them cached on each machine. In our method, we adopt the mechanism [4] dealing with join operation, which broadcasted the schema triples to every computing node before the join operation. For the rules containing multi-join operation, the join operation inside the schema triples can

Algorithm 1 OWL reasoning algorithm on Spark

```
Input: triples, OWL Horst rule set
Output: result
schema_derived=triples.apply(schema_rules)
val flag=true
while flag do
    SPO_derived=triples.apply(SOP_rules)
    if SPO_derived!=null then
        triples=triples.union(SPO_derived)
    end if
    type_derived=triples.apply(type_rules)
    if type_derived!=null then
        triples=triples.union(type_derived)
    end if
    if SPO_derived==null && type_derived==null then
        flag=false
    end if
end while
sameAs_derived=triples.apply(sameAs_rules)
triples=triples.union(sameAs_derived)
return triples
```

firstly be executed in memory locally. The instance triples will be divided into many partitions and then execute the join operation between schema triples and instance triples in parallel. Here we take the O16 as an example to describe the operation. The execution procedure is described in Algorithm 2.

Algorithm 2 Optimized the join operation of rule O16 using broadcast variables

```
Input: triples
Output: results
val triples=sc.textFile("hdfs://...")
val op=triples.filter(t⇒t._2.equals("owl:onProperty")).map(t⇒(t._1, t._3))
   .collect.toMap
val opBroadcast=sc.broadcast(op)
val av=triples.filter(t⇒t._2.equals("owl:allValuesFrom")).map(t⇒(t._1, t._3))
   .collect.toMap
val avBroadcast=sc.broadcast(av)
val ins=triples.filter(t⇒av.value.contains(t._1) && op.value.contains(t._1))
   .map(t⇒((t._1, op.value(t._2)), av.value(t._2)))
val typ=triples.filter(t⇒op2.contains(t._2) && av.value.contains(op2(t._2)))
   .map(t ⇒((t._1, t._2), t._3))
val results=ins.join(type.map(t⇒(t._2._2, t._2._1)))
return results
```

Reasoning for *owl:sameAs* property The sameAs rules derived too many triples and most of the output data is valueless in practical applications [18]. To improve the performance, many existing methods prefer to build a sameAsTable [18], [4], which design a hash-code function and compare the hash value. However, dealing with *owl:sameAs* property often consume too much time. Based on the character of dataset, we do not adopt the above method but simplify the reasoning process. The O6 can also be dealt with the transitive closure Algorithm 3.

Algorithm 3 transitive closure algorithm

```
Input: transitive triples
Output: results
var flag=1L
var p=triples
var q=p
var l=q.map(t ⇒((t._2, t._3), t._1)).partitionBy(partitioner)
while flag!= 0 do
    val r=q.map(t ⇒((t._2, t._1), t._3)).partitionBy(partitioner)
    val q1=l.join(r).map(t ⇒(t._2._1, t._1._1, t._2._2))
    q=q1.subtract(p, parallism).persist(storagelevel)
    flag=q.count()
    if flag!=0 then
        l= q.map(t ⇒((t._2, t._3), t._1)).partitionBy(partitioner)
        val s=p.map(t ⇒((t._2, t._1), t._3)).partitionBy(partitioner)
        val p1=s.join(l).map(t ⇒(t._2._2, t._1._1, t._2._1)).persist(storagelevel)
        p=p1.union(q).union(p)
    end if
end while
return p
```

5 EVALUATION

In this section, we conduct a series of experiments to compare the performance of the proposed approach with the KP [6] and Cichlid [4] under the same environment.

5.1 Experiment and Dataset

We set up a cluster with one master and four worker nodes. Each node has 48 Xeon E5_4607 2.20GHz processors, 64GB memory and 10TB 7200 RPM SATA hard disk. The nodes are connected with Gigabit Ethernet. All the nodes run on 64-bit Ubuntu 12.04 LTS operating system and Ext3 file system. The version of the Spark is 1.0.2. And the corresponding Hadoop v2.2 with Java1.7 is installed on this cluster. The Spark has special requirement for the version of Hadoop. Besides, the version of the Scala is 2.10.6. We use the synthetic benchmarks. LUBM [5] is a widely used standard benchmark for semantic program. Due to the limitation of hardware, we use the data generator UBA to generate 5 sets of data with different universities: LUBM_10, LUBM_50, LUBM_100, LUBM_150 and LUBM_200 in our experiment. The number of triples in each data set is shown in Table 4 .

Table 4: The number of triples for each dataset

dataset	LUBM_10	LUBM_50	LUBM_100	LUBM_150	LUBM_200
triples	1449832	6822632	12297500	20652819	27643644

5.2 Experiment Performance

We evaluate the performance of our method with KP and Cichlid, where KP adopts the executable strategy in [6]. All experiments run three times and the average value is listed as follow.

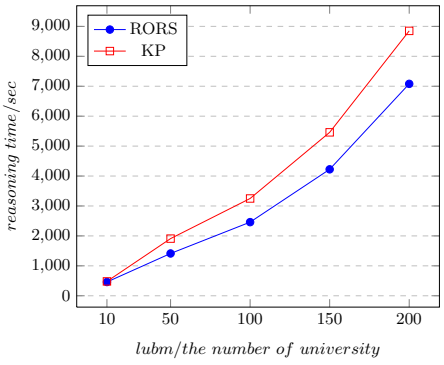

Fig. 6: The runtime of RORS and KP

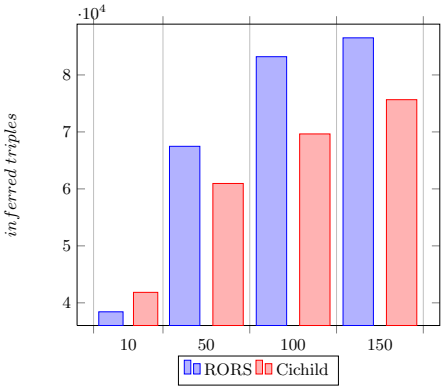

Fig. 7: The number of triples per sec

Figure 6 displays the reasoning time of our approach and KP with different scale of data sets. The reasoning time includes the time of dividing input data and eliminating duplicated triples. We can see that the reasoning time of RORS and KP increase almost linearly with the growth of data size. The result shows that our approach is better than the KP under the same environment. The performance of reasoning is improved by 30% approximately.

Figure 7 shows the number of inferred triples per second. Our method can infer more implicit triples than Cichlid and the performance is improved by 26% approximately.

6 Discussions

In this paper, we present an approach to enhancing the performance of the rule-based OWL reasoning on Spark based on a locally optimal executable strategy. Our method performs better than KP in reasoning strategy. Although the approach does't find the optimal executable strategy for global rules, our method can be used as the valuable foundation for future research in the rule-based OWL reasoning. Therefore, in the future work, we plan to design some algorithms to find the optimal strategy for global rules.

There are many works to develop OWL reasoning systems, including early works, such as Pellet [14], Jena [9], and Sesame [1], etc. These reasoners use a composition tree model for ontology to infer implicit information and exhibit both large time and space complexity. Due to the limitation of computing resource and running speed, these systems can hardly achieve excellent performance. Therefore, many distributed reasoning system emerged. In [10] and [12],

they proposed a parallel reasoning method that the reasoning rules are executed repeatedly until no extra data is generated. But there exists much more data communication cost. In [20], Weaver and Handler proposed a data partitioning model based on MPI, but this method do not filter duplicate data. [17] presented a distributed reasoning system which based on MapReduce. It analyzed the dependency between rules and builded a dependence graph. But it generated large amount of useless middle data and huge data communication cost. Then Urbani proposed a MapReduce-based parallel reasoning system with OWL-Horst rules called WebPIE [18]. It can deal with large scale ontology on a distributed computing cluster. However, WebPIE exhibits poor reasoning time. [16] are rule-based OWL reasoner. Although they can infer large scale triples, it costs too much reasoning time.

[13] proposed a rule-based reasoner that used massively parallel hardware to derive new facts based on a given set of rules, but that implementation was limited by the size of processable input data as well as on the number of used parallel hardware devices. Seitz [15] presented an OWL reasoner for embedded devices based on CLIPS. Urbani [19] proposed a hybrid rule-based reasoning method that combined forward and backward chaining, and implemented a prototype named QueryPIE. Terminological triples are pre-computed before query, which is used to speed up backward-chaining at query time. In [4], although the author has improved the reasoning time using Spark, it ignores the analysis of interdependence among the rules and does not give the optimal executable strategy. Besides, MPPIE [8] recently implemented the RDFS reasoning on Giraph.

7 Acknowledgments

This work is supported by the program of the National Natural Science Foundation of China (NSFC) under 61502336, 61373035 and the National High-tech R&D Program of China (863 Program) under 2013AA013204.

References

1. Broekstra J., Kampman A., & Van Harmelen F. (2002) Sesame: A generic architecture for storing and querying RDF and RDF schema. In: *Proc. of ISWC 2002*, Springer, pp. 54–68.
2. Baader F. & Sattler U.(2001) An overview of tableau algorithms for Description Logics. *Studia Logica*, 69(1):5–40.
3. Chang L., Guilin Q., Haofen W. & Yong Y. (2011) Large scale fuzzy pD* reasoning using MapReduce. In: *Proc. of ISWC 2011*, Springer, pp. 405–420.
4. Gu R., Wang S., Wang F., Yuan C., & Huang Y.(2015) Cichlid: Efficient large scale RDFS/OWL reasoning with Spark. In: *Proc. of IPDPS 2015*, IEEE, pp. 700–709.
5. Guo Y., Pan Z., & Heflin J.(2005) LUBM: A benchmark for OWL knowledge base systems. *J. Web Sem.*, 3(2-3):158–182.
6. Kim J. & Park Y. (2015) Scalable OWL-Horst ontology reasoning using Spark. In: *Proc. of BigComp 2015*, pp. 79–86.

7. L. McGuinness D., van Harmelen F. (2004) Web Ontology Language. *W3C Recommendation*.
8. Lv X., Wang X., Feng Z., Rao G., Zhang X., & Xu G.(2016) MPPIE: RDFS parallel inference framework based on message passing (*in Chinese*). *Journal of Frontiers of Computer Science and Technology*, 10(4): 451-465.
9. McBride B. (2002) Jena: A semantic web toolkit. *Internet computing 2002*, IEEE, 6(5):55–59.
10. Mühleisen H. & Dentler K. (2012) Large-scale storage and reasoning for semantic data using swarms. *IEEE Comp. Int. Mag.*, 7(2):32–44.
11. Maeda R., Ohta N., & Kuwabara K.(2014) MapReduce-based implementation of a rule system. In: *Recent Developments in Computational Collective Intelligence 2014*, Springer, pp. 197–206.
12. Oren E., Kotoulas S., Anadiotis G., Siebes R., ten Teije A., & van Harmelen F.(2009) Marvin: Distributed reasoning over large-scale semantic web data. *J. Web Sem.*, 7(4):305–316.
13. Peters M., Brink C., Sachweh S., & Zündorf A. (2013) Rule-based reasoning on massively parallel hardware. In: *Proc. of SSWS 2013 at ISWC*, pp. 33–49.
14. Sirin E., Parsia B., C. Grau B., Kalyanpur A., & Katz Y. (2007) Pellet: A practical OWL-DL reasoner. *J. Web Sem.*, 5(2):51–53.
15. Seitz C., & Schönfelder R. (2011) Rule-Based OWL reasoning for specific embedded devices. In: *Proc. of ISWC 2011*, Springer, pp. 237-252.
16. Thakker D., Osman T., Gohil S., & Lakin P.(2010) A pragmatic approach to semantic repositories benchmarking. In: *Proc. of ESWC 2010*, Springer, pp. 379–393.
17. Urbani J., Kotoulas S., Oren E., & Van Harmelen F. (2009) Scalable distributed reasoning using mapreduce. In: *Proc. of ISWC 2009*, Springer, pp. 634–649.
18. Urbani J., Kotoulas S., Maassen J., Van Harmelen F., & Bal H.(2010) OWL reasoning with WebPIE: Calculating the closure of 100 billion triples. In: *Proc. of ESWC 2010*, Springer, vol.1, pp. 213–227.
19. Urbani J., van Harmelen F., Schlobach S., & E. Bal H. (2011) QueryPIE: Backward reasoning for OWL Horst over very large knowledge bases. In: *Proc. of ISWC 2011*, Springer, pp. 730–745.
20. Weaver J. & A. Hendler J. (2009) Parallel materialization of the finite RDFS closure for hundreds of millions of triples. In: *Proc. of ISWC 2009*, Springer, pp. 682–697.
21. Wu H., Liu J., Ye D., Zhong H., & Wei J. (2013) A distributed rule execution mechanism based on mapreduce in sematic web reasoning. In: *Proc. of Internetware 2013*, ACM, Article No. 6.
22. Zaharia M., Chowdhury M., J. Franklin M., Shenker S., & Stoica I. (2010) Spark: Cluster computing with working sets. In: *Proc. of HotCloud 2010*, Boston, MA, USA.
23. Zaharia M., Chowdhury M., Das T., Dave A., Ma J., McCauley M., J. Franklin M., Shenker S., & Stoica I. (2012) Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In: *Proc. of NSDI 2012 at USENIX*, pp. 15–28.